Applying Decision Analysis to PATH: Discussion Paper

C. Peters
D. Marmorek

1.0 Background

Decision analysis is a method used to facilitate decision-making by organizing all relevant information about management actions and objectives, performance measures, key hypotheses, and uncertainties about these hypotheses into a systematic framework. The method was first developed in business management science, but has since been employed in fisheries and other resource management problems.

The idea of applying decision analysis to PATH was first proposed by Saul Saila and Carl Walters in the SRP review of the Preliminary Report on Retrospective Analysis in April of 1996 as a way to evaluate management options and alternative experimental designs. These ideas were expanded on at the April 1996 PATH workshop at Kah-Nee-Tah by Randall Peterman of SFU. Further sub-group discussions at the workshop resulted in a preliminary decision tree showing the major components of a Columbia River decision analysis, and various methods for assigning probabilities to alternative hypotheses. The need for a structured approach to decision making was reiterated in the SRP review of the Final Report on Retrospective Analysis in September 1996, and the concept was again discussed in plenary and sub-group meetings at the Wenatchee workshop in October 1996. Dr. Peterman also presented a short primer on decision analysis, including some examples of its application to resource management decision, at this workshop.

Development of a decision analysis framework for spring/summer chinook was identified as a prospective task for FY97 at the last PATH workshop in Wenatchee. As a first step in this process we are circulating this discussion document, which lays out our preliminary thinking on the structure and implementation of decision analysis as applied to PATH and Columbia River salmon. In the document, we propose some alternative ways of approaching some of the major issues and discuss some of the advantages and disadvantages of these alternatives. Our purposes are to:

- provide something for PATH members to think about;
- elicit feedback on the feasibility of alternative approaches (i.e. availability of data, appropriateness of required analyses); and
- ensure that the development of the structure of the decision analysis is consistent with ongoing PATH analyses of specific components, particularly the decision tree work in Chapter 6 and the prospective modelling work.

The first section of the document contains general discussion of some of the issues presented at the Wenatchee workshop. Our intention is to provide some alternative approaches to dealing with these issues; we do not expect that all of these issues will necessarily be resolved right away. In the second section, we describe a preliminary example of a decision analysis for spring/summer chinook. The purpose of the example analysis is to further demonstrate, using actual numerical examples, how decision analysis tools might be implemented in PATH. A more detailed pilot decision analysis will be completed in early 1997 (as outlined in the FY97 Work Plan). Robin Gregory, an expert in decision analysis and elicitation of expert opinion, will assist the PATH group in the preparation of the pilot analysis. Dr. Gregory was recommended to us by Randall Peterman, and has already given us some valuable suggestions during two brief meetings.

The process of developing a decision analysis for PATH should be viewed as an experiment. The most important product of the analysis will be insight into how differences between key hypotheses, uncertainties, evidence, and assumptions affect management decisions. In addition, we emphasize that decision analysis is simply a way to organize existing information (including quantitative models, empirical data, and explicit scientific assumptions and judgements) into a systematic and structured framework for decision-making. Therefore, it should be seen as complementary to other decision-making tools, such as the hierarchical decision tree developed in Ch. 6. In fact, we convert Ch. 6 into a decision tree amenable to decision analysis methods in Appendix A of this document to show the complementary nature of these two approaches.

For the purposes of this document, it is convenient to divide the decision analysis into two separate analyses. The first is an analysis of actions that are designed to bring about recovery and survival of listed chinook stocks. The second is a decision analysis of experimental management actions that are intended to enhance learning and resolve uncertainties about alternative hypotheses. In this document, we discuss only those issues relevant to the analysis of recovery actions. Although the ultimate intention is to pursue both analyses, we chose to focus only on the analysis of recovery actions for now because the concepts and methods can be illustrated best with a very simple example. Also, the decision analysis of recovery actions will provide useful insights to generate ideas on experimental management actions. Since the purpose of decision analysis is to clarify the decision problem, it is best not to overwhelm ourselves with complexity initially. Eventually, it may be desirable to apply the method to a larger set of actions, stocks, and performance measures (e.g. performance measures and strategies implied by the Return to the River report).

2.0 General Discussion of Issues in Applying Decision Analysis

One of the features of decision analysis is that it takes a potentially complex decision problem and breaks it down into its smaller and more manageable components. Once analyzed, the individual components can then be reassembled into a systematic framework for making decisions. There are essentially 8 general components of decision problems, which have been described by Randall Peterman at the last couple of PATH Workshops. The eight essential elements of decision analysis include:

1. a list of alternative management actions (i.e. different experimental management plans);
2. management objectives composed of performance measures (e.g. ln(R/S)), which are criteria for ranking management actions;
3. uncertain states of nature (e.g. different hypotheses about key relationships such as flow survival);
4. probabilities of those states;
5. model to calculate the outcomes of each combination of each management action and each hypothesized state of nature;
6. decision tree;
7. rank actions based on the expected value of the performance measures; and
8. sensitivity analyses.

Discussion of the issues relating to PATH are organized into these 8 components.

2.1 Alternative Management Actions

What management actions should be included/excluded on the basis of the retrospective results?

Rather than excluding management actions on the basis of retrospective analyses, it may be better to simply prioritize management actions on the basis of how much improvement is possible. For example, if habitat and hydro actions were expected to provide the most opportunity for recovery and survival of spring-summer chinook, initial analyses should focus on these actions.

Should packages of actions include hydro, habitat, hatchery and harvest components or should we do separate analyses on each H first?

Preliminary analyses should focus on individual actions until the technical details of the analysis have been worked out. However, actions within the different H's should be considered together in packages as soon as technically possible since this most closely resembles the actual decision-making approach taken by the region. Analysis of combinations of actions may also help to identify new actions and strategies that would otherwise not be apparent. These new insights can then be used to further refine analyses of individual actions.

Using packages of actions presents two problems. First, the number of possible combinations of different management actions increases the complexity of the analysis. Second, it will require some assumptions to be made about the interactions between different actions. For example, a particular habitat action is likely to be more effective in improving survival when combined with some reduction in passage mortality than when the habitat action is implemented alone. Similarly, actions to improve the **quality** of habitat will presumable make the greatest difference at low population, low survival conditions, whereas actions to improve the **quantity** of habitat may be more important at moderate population sizes. The direction (i.e. synergistic vs. antagonistic effects) and magnitude of these interactions will have to be considered when evaluating suites of actions from the different H's.

Who should be involved in recommending packages (IT? which committee?)

Perhaps the PATH subcommittee of the IT could respond to the general structure of the pilot analysis at the PATH briefing scheduled for February.

2.2 Performance Measures

What biological performance measures should be used?

Performance measures are quantitative indicators of management objectives. Some performance measures for survival and recovery standards have already been set by the NMFS jeopardy standards (i.e. probability that the number of spawners will exceed defined survival and recovery thresholds). Comparing these probabilities for different management actions can be used to show which actions are most effective in attaining these goals (i.e. maximizes the probability of exceeding the threshold). Other decision criteria could also be applied here, such as a maxi-min criterion (choose the action whose worst-case outcome or worst quartile of outcomes is the highest; a risk-averse criterion).

The decision analysis should be designed to compute other performance measures in addition to the NMFS jeopardy standards. For example, probabilities can be calculated for a range of different thresholds to show how the preferred management action depends on the threshold level. Alternatively, a frequency distribution of the numbers of spawners under proposed management actions could be generated to show these probabilities graphically (such as those shown in Figure 1.A.1 of the BRWG 10/13/94 Progress Report). The main point here is that the framework must be flexible to allow the use of a number of possible performance measures. The pilot decision analysis will demonstrate some of these alternative forms of performance measures.

Other biological performance measures besides survival and recovery may also be desired. For example, the System Configuration Team and several Army Corps of Engineers reports have evaluated the effects of various changes to the hydrosystem. In addition, the ISG report "Return to the River" identifies other biological performance measures, such as the life history diversity of salmon populations or the interconnectedness, diversity, and complexity of fish habitat. The PATH decision analysis should rely on these other reports and analyses as much as possible to identify performance measures that are not directly related to the NMFS jeopardy standards. This suggests that management actions evaluated in the decision analysis should also be consistent with actions considered in these other reports. Techniques are available to construct quantitative indices for qualitative performance measures (e.g. diversity of life histories) to allow explicit trade-offs across performance measures. If these techniques are not used, a more qualitative form of evaluation may be necessary. Sensitivity analyses (section 8.) will be especially important for these qualitative judgements.

Are there other non-biological performance measures (e.g. social and economic impacts) that should be considered?

Although the principal focus of the analyses to date has been on the biological implications of proposed recovery actions, it is reasonable to expect that social and economic considerations will eventually play some role in decision making. Identifying these non-biological implications are clearly outside the scope and expertise of PATH. However, economists from NMFS and the NPPC are working on identifying economic issues in recovery actions (the NPPC meeting in Kelso, Wa last week discussed the appointment of an Independent Economic Analysis Board). Another potentially important performance measure to include is the degree of learning over time associated with a particular management action. This is especially relevant to adaptive management actions, but may apply to survival/recovery actions as well. The decision analysis framework should be designed to accommodate these kinds of performance measures if and when they become necessary.

Will different performance measures be combined into a single criterion to rank management actions or be kept separate?

Although there are quantitative methods (e.g. multi-attribute utility theory) that could be used to combine different performance measures into a single value, the number and diversity of performance measures and interest groups in the Columbia River virtually guarantees that these methods will be expensive, time-consuming, and ultimately unsuccessful. An alternative to these more formal quantitative methods is to elicit the relative importance of one performance measure over another from individuals or groups. This could be done for several different stakeholder groups on an experimental basis, possibly using the PATH subcommittee of the Implementation Team as a partial microcosm of the diversity of stakeholders. This would help to determine the sensitivity of preferred decisions to different trade-off assessments. Robin Gregory has considerable experience in doing these elicitations, and has recently employed such an approach in an analysis of alternative sites for a nuclear waste depository.

The other alternative is a multiple accounts approach. In this method, different performance measures calculated for a given management action are presented to decision-makers together in a table, allowing them to make their own trade-offs. A multiple accounts approach avoids some of the assumptions and technical demands that would be required to combine multiple performance measures, but places more demands on the decision-maker.

Will decisions be made on individual stocks or basin-wide ?

It makes some sense to treat stocks individually in the decision analysis because management objectives (e.g. BRWG recovery targets), data availability, and some management actions (particularly habitat and hatchery actions) are likely to be stock-specific. Hydro and harvest decisions would apply to larger stock groupings, and should be assessed over a number of stocks and species.

The disadvantage of basing decisions on individual stocks is that it is not clear to what degree the index stocks reflect the response of the entire ESU to hatchery and habitat actions. Therefore, decisions made for one stock may not necessarily be directly applicable to other stocks in the region. Work by the habitat group to rate habitat using the Eastside Assessment GIS may be helpful here in developing region-wide assessments of habitat effects.

Can the decision analysis be expanded to include other species (fall chinook, steelhead, sockeye)?

Since similar decisions will eventually have to be made for other stocks and species (particularly fall chinook), the decision analysis framework should be flexible enough to consider these species in the future. However, to keep the analysis as simple as possible to start with, the decision analysis should look only at spring/summer chinook before extending the approach to others. A more qualitative analysis may be necessary for some species where data is lacking. In these cases, it is probably better to keep the analysis explicitly qualitative rather than pseudo-quantitative.

2.3 Uncertain States of Nature

Is it feasible to develop alternative hypotheses by life stage, or is it better to look at impacts on overall life-cycle survival?

Early discussions of decision analysis at Kah-Nee-Tah focussed on developing hypotheses about the effects of different management actions on individual life stages, then progressing through the life stages to project overall survival rates and spawner abundances. There are three significant challenges with this approach:

1. Responses to a given management action will in many cases depend on the response in a previous life stage. For example, the response to reductions in passage mortality from some hydro action may depend on habitat conditions during freshwater life stages. Therefore, probabilities placed on uncertain responses to management actions will be conditional on responses to previous life stages. This greatly complicates the procedure for deriving these probabilities, and increases the computational demands of the analysis.

2. Data are only available for some life-history stages (e.g. FSR, JMC) but not others (e.g estuarine/ocean). Furthermore, FSR survival estimates (Ch. 9) are for the aggregate of stocks, not individual index stocks.

3. A life-stage approach is not consistent with the Bayesian prospective model, which is based on Stock/Recruit relationships.

There are some complementary approaches which might achieve the best of both worlds. The first is to use nested models to develop aggregate hypotheses from life-stage specific responses to management actions. For example, actions to improve habitat could be represented as increases in FSR survival rates. These survival rates could then be used in a life cycle model to generate estimates of the change in spawner abundance resulting from this change in FSR survival. Such aggregate hypotheses could be incorporated into the Bayesian prospective model if they were expressed in terms of stock-recruit parameters. Aggregate hypotheses developed using nested models would have to be constrained by observed temporal and spatial patterns in overall survival, as well as life-stage specific survival. In addition, the prospective model and life-cycle models could be used iteratively to test the internal consistency of their predictions.

An aggregated life-cycle approach has some advantages over the life-stage specific approach, but also has its own shortcomings. In an aggregated life-cycle model (like the current prospective model), one would directly represent predicted effects of management actions as changes in stock-recruitment parameters. Under this approach, a given habitat action might be expressed as some increase in the Ricker 'a' parameter. Reductions in harvest rates would be represented as reductions in instantaneous mortality rates. Similarly, alternative versions of the passage models (tested against PIT-tag and other survival studies, transportation studies, and MLE estimates of μ to derive posterior probabilities, as discussed in the Kah-Nee-Tah report) could be used to estimate changes in μ in an overall life-cycle model under different actions.

Whatever approach is used, expert opinion can be used to make the links between management actions and overall survival when data is not available. However, the difficulty with this approach is that the link between specific management actions and the a, μ, and other parameters in a Ricker stock-recruitment function may be abstract. Therefore, it may be difficult to explain the implications of parameter choices when eliciting expert opinions. These difficulties may be overcome through an iterative procedure that reveals consequences of different choices.

What are some possible uncertainties to consider?

For logistical reasons, it is usually necessary to limit the number of uncertain states of nature. The decision analysis should therefore focus on those uncertainties that have the greatest effect on the outcomes of management actions. Retrospective analyses, such as the decision tree in Ch. 6, will help to determine these key uncertainties. Possibilities include:

- effects of actions to improve habitat quality/quantity on overall life-cycle survival;
- effects of hatchery actions on various life stages and overall life-cycle survival;
- effects of hydro actions (particularly transportation and natural river drawdown) on JMC survival and overall life-cycle survival (including delayed effects);
- variability in estuarine/ocean survival;
- effects of hydro and harvest actions on up-river passage survival and overall life-cycle survival; and
- uncertainties in stock-recruitment curves, including the existence and magnitude of depensatory mechanisms and the shape of the curve at high spawner densities (i.e. Ricker-shaped vs. Beverton-Holt)

2.4 Calculate Probabilities for Uncertain States of Nature

How will probabilities be assigned to alternative hypotheses / uncertain states of nature?

If data exists, MLE/Bayesian methods or quantitative comparison of model output to different data sets can be used to assign probabilities (discussed in Kah-Nee-Tah workshop report, pp. 37-38). For JMC/hydro actions, for example, one could compare different hypotheses about direct and delayed survival with several data sets (e.g. stock-recruit data, PIT-tags, NMFS survival studies, transportation studies). Similarly, for FSR habitat models, one could generate different stock-recruit curves given a more detailed representation of effects of management actions on life history, which could then be used to generate a probability distribution of Ricker ‘a’ and ‘b’ values. Alternatively, probabilities may be generated by comparing Ricker ‘a’ values for stocks in different qualities of habitat. In either case, some expert judgement will probably be required to assign probabilities.

Using expert opinion to generate probabilities for uncertain states of nature is commonly used in developing decision analyses, and there are various techniques available to do this. Robin Gregory has expressed interest in working with us to elicit these probabilities from PATH members and implementing a decision analysis framework. We propose that Robin conduct a series of short meetings in early 1997 to generate these probability distributions, and assess as a group how this process works.

Several concerns have been raised about using expert opinion to quantify uncertainties:

- it is less rigorous than basing these estimates on empirical data;
- the approach is not consistent with PATH’s emphasis on using empirical data to resolve long-standing disagreements over alternative hypotheses; and
- conflicting opinions by different experts may revive previous arguments over alternative hypotheses and prevent any progress from being made.

These are valid concerns. Using empirical data to estimate the effects of management actions and their uncertainties is the preferred approach (providing that the methods used are sound) and should be used whenever possible. However, there will almost certainly be cases where empirical data is either unavailable or insufficient to allow a thorough analysis. In these cases, scientists can give decision-makers three kinds of advice:

1. They can recommend that a decision should not be made until more data is available. Even if their recommendation specifically identifies what data is needed and when it will be available, this advice does not address the short-term need for information when a decision must be made. In addition, it should be recognized that not making or postponing a decision is in itself a decision (i.e. a decision to maintain the status quo), which can have serious consequences.

2. They can tell the decision-makers that they don’t know what the effects of all possible management actions might be, and leave it to decision-makers to decide what to do. In this case, decision-makers may use their own judgement to make the decision, based on what they know or have been told.

3. Scientists or other experts can provide their best judgement about what the effects might be along with the estimated uncertainties in these judgements. This approach is preferable to 1 and 2 because a) it allows decisions that cannot be delayed to be made based on all relevant information, including indirect evidence, informal observations, and the accumulated experience of experts familiar with the system; and b) in most cases, the judgement of experts who are most familiar with the alternative

hypotheses and related evidence will be better than that of decision-makers that are less familiar with these hypotheses.

The third concern listed above is also a potential problem, at least to the extent that different opinions about alternative hypotheses lead to different recommended management actions. Therefore, it will be important to test the sensitivity of recommended decisions to different judgements about hypothesized effects of management actions and their uncertainties. These differences could result from basing expert opinions on different kinds of information, as well as from soliciting the judgements of different experts. If the decisions are insensitive to the judgements, then differences in opinion are of no consequence. If there are differences, such analyses can help to identify why these differences exist.

Important evidence may become available between the time when a decision analysis framework is constructed and when a decision has to be made. For example, estimates of recruits/spawner for several brood years in the early 90s when flows were high will become available in the next couple of years. This should be kept in mind when constructing the framework so that any new or updated information can be accommodated as easily as possible.

2.5 Model to Calculate Outcomes

What model or models should be used to simulate the outcomes of alternative actions?

The model used to project the outcomes of different management actions is the focus of current work in the prospective arena. Prospective analyses are intended to identify what improvements in overall survival (and thus in the number of spawners) of Snake River spring/summer chinook stocks are possible from various management actions within each of the 4 H's (Habitat, Harvest, Hatchery, and Hydro). There are currently 4 life-cycle models and 3 passage models that are available to simulate the outcomes of these actions. The life-cycle models include: the Bayesian population model being developed by Deriso; SLCM; ELCM; and SPM. The 3 passage models are CRiSP, FLUSH, and PAM. Which of these models will be used in prospective analyses or how different models might be combined has yet to be finalized. Our initial thinking is to use the Bayesian life-cycle model to do the actual simulations, but to incorporate information from a subset of the other models to affect the range of parameter values used in the Bayesian model.

An important consideration in applying any of these models in a prospective context is to identify specific parameters in the models that can be adjusted to represent the effects of management actions. For example, the outcomes of actions to improve FSR habitat quality might be simulated by increasing egg to smolt survival rates. An action that is expected to produce only small increases in habitat quality will be associated with a small increase in egg-smolt survival rates, while an action that is expected to produce large increases in habitat can be represented by a larger increase in egg-smolt survival. By equating different management actions with different values of these parameters, the relative performance of different actions can be quantified by running the model using different parameter values and comparing the outcomes. This can be done in a decision analysis framework by representing a management action with several different parameter values rather than just one and comparing the relative performance of different actions under a range of possible outcomes.

As a first step towards doing prospective and decision analyses, we have identified a number of potential "management" variables for each of the four H's (see Table 1). In general, habitat actions are expected to increase overall productivity and/or freshwater survival and rearing capacity. Harvest actions can be simulated by adjusting simulated harvest rates. Effects of hatchery actions are more difficult to assess, but may affect overall productivity if the survival of hatchery fish is worse than wild fish or if release of hatchery fish leads to density-dependent reductions in survival. For some models, hatchery effects can be modelled directly by assigning different survival rates to wild and hatchery fish or by modelling interactions between these two groups. Hydro actions can be modelled either by:

1. running different scenarios in passage models, assigning a posterior probability to each run based on the past performance of the model (as described in the Kah-Nee-Tah workshop report), and inputting this information into life cycle models; and

2. estimating the response of passage survival rates to different hydro management actions directly (e.g. by synthesizing studies of dam passage and SAR survival, as was done in Ch. 6).

Table 1. Model parameters that could be used to simulate management actions.

	Habitat	**Hatchery**	**Harvest**	**Hydro**	**Comments**
Bayesian model $\ln(R)=a + \delta\text{-}m + \ln(\beta)$ $+ (1+p) \ln(S)$ $- (1/\gamma) \ln(1$ $+\gamma\beta S)$	S-R parameters: a b ($b=\ln(\beta)$) σ^2 of a	S-R "a" parameter density - dependent effects of ΔS_h	Simulated in-river harvest rate (either from step function or actual vs intended harvest rate rel.)	m (inriver survival), either input from passage models or estimated from other sources	
Information Sources	Subjective habitat rankings in MLE analysis (planned by habitat group) Ch. 9, Ch. 4 analyses Expert opinion	Contrasts in productivity or survival rates across streams (Ch. 11) Expert Opinion	Harvest rate data (Ch. 13) Expert Opinion	Ch. 5 and 6 data and analyses Expert Opinion	
Other sources of information that can feed into the Bayesian model					
Life-Cycle Models	Parameters in survival rate relationships for freshwater life stages (egg-presmolt, pre-smolt-smolt, in-basin smolt survival) Carrying capacity in freshwater habitat	Density-dependent effects in FSR habitat Differential survival and productivity of hatchery fish relative to wild at diff. life stages Interactions between wild/hatchery fish			
Passage Models				Passage survival rates	

We have also included a row for information sources that could be used to estimate the response of model parameters to different management actions.

This list is very preliminary and is intended only to provoke some thought about how each model could be used for prospective or decision analyses. We expect that those that are more familiar with each of the models will be able to modify the list, provide comments, and identify data sets that might provide insight into how to proceed with prospective and decision analyses.

A major issue in evaluating the effects of different actions is how to combine different effects. For example, how can a change resulting from some habitat action be evaluated concurrently with a change in μ resulting from a hydro action? Are these effects additive, or are there positive/negative effects associated with implementing them both? What is an appropriate mathematical structure for modelling multiple effects? The value of integrating all actions into a relatively simple life-cycle model is that the net cumulative benefit of a number of actions can be realistically constrained, and easily understood.

2.6 Decision Tree

The final structure of the decision tree will depend on the resolution of the issues and alternatives described above. In the example decision analysis, we consider a structure without separate life history stages to begin with (see Section 3 for more details). The decision trees created from Ch. 6 information in Appendix A also aggregates over all life history stages between smolts and recruits. However, these analyses can be modified as more information is acquired.

2.7 Rank Actions

It will very likely not be possible to get a single answer out of the decision analysis that is satisfactory to everyone because disagreements over assumptions, objectives, methods, etc. are bound to occur. However, the process of looking at these issues in a decision-analysis context will provide some common ground for discussion, distinguish critical differences of opinion from ones that do not significantly affect the ranking of decisions, and help to clarify the effects of major uncertainties and assumptions. This will in turn provide guidance to the detailed design of adaptive management experiments to address critical uncertainties, as well as helping to further refine ongoing retrospective analyses.

2.8 Sensitivity Analyses

The complexity of the biological and management system virtually guarantees that not everyone will agree with certain assumptions, approaches, or parameter values. Therefore, sensitivity analyses will be an essential component of the decision analysis to identify those points of contention that most affect the decision. Example sensitivity analyses include:

- weightings of different data sets used to estimate posterior probabilities;
- subjectively assigned probabilities and probability distributions (e.g. change in ‘a’ value resulting from habitat actions); and